arXiv:1604.07242v1 [cs.MS] 25 Apr 2016

Implementation of *hp*-adaptive discontinuous finite element methods in Dune-Fem

Christoph Gersbacher

Department of Applied Mathematics, University of Freiburg,
Hermann-Herder-Str. 10, D-79104 Freiburg, Germany.

April, 2016

Abstract

In this paper we describe generic algorithms and data structures for the implementation of *hp*-adaptive discontinuous finite element methods in the Dune-Fem library. Special attention is given to the often tedious and error-prone task of transferring user data during adaptation. Simultaneously, we generalize the approach to the restriction and prolongation of data currently implemented in Dune-Fem to the case of *p*- and *hp*-adaptation. The `dune-fem-hpdg` module described in this paper provides an extensible reference implementation of *hp*-adaptive discontinuous discrete function spaces. We give details on its implementation and the extended adaptive interface. As proof of concept we present the practical realization of an *hp*-adaptive interior penalty method for elliptic problems.

1 Introduction

Adaptive mesh refinement or *h*-refinement is today a standard tool in the acceleration of finite element methods. The *hp*-version of the finite element method was made popular by Babuška and co-workers in a series of papers [2, 16, 17, 18, 19, 20]. They showed that in case the solution is sufficiently smooth an adaptation in p, i.e., in the polynomial order of a local approximation space, may be more advantageous. By a proper combination of *h*- and *p*-refinement exponential convergence towards the exact solution in terms of the number of degrees of freedom may be observed.

From a theoretical point of view *hp*-adaptive finite element methods are a natural generalization of *h*-adaptive methods. However, their practical implementation still poses considerable difficulties and few open source software solutions are available to researchers. Of these we want to mention the C++libraries *Concepts*[26, 15] and *deal.II*[3, 4], and the Fortran 90 library *PHAML*[27].

In the present paper we want to discuss data structures and algorithms for implementing *hp*-adaptive discontinuous finite element methods in the Dune-Fem finite element library [10]. The main difficulty we want to address is the restriction and prolongation of user data in adaptive computations. In case of local mesh refinement utility classes in Dune-Fem relive users of this often tedious and error-prone task. The dedicated data structures and algorithms need to be generalized to the case of *p*- and *hp*-refinement in such a way that legacy code remains valid.

At the time of writing there has been one prior attempt to implementing *hp*-adaptive finite element methods in DUNE-FEM. In a pioneering work Andreas Dedner and Robert Klöfkorn explored how to implement continuous and discontinuous *hp*-adaptive finite element spaces in DUNE-FEM. For the most part, the restriction and prolongation of discrete functions may be considered as beyond the scope of this implementation and was done in a preliminary way only. For the local approximation Lagrange ansatz polynomials are used. The software was used in the preparation of numerical results for stationary problems in [9].

This paper is accompanied by a new extension module to the DUNE-FEM library. The `dune-fem-hpdg` module provides a number of discrete function spaces for implementing *p*- and *hp*-adaptive discontinuous finite element methods in DUNE-FEM. Hardly any restrictions to the local ansatz functions are imposed. The software is extensible and allows users to quickly setup new finite element spaces using customized ansatz functions with minimal effort. The restriction and prolongation of discrete functions and other user data is handled in a fully automated fashion. The software is compatible with DUNE-FEM 2.4 and in particular with its data structures and algorithms for local mesh refinement.

The outline of this paper is as follows. In Section 2 we briefly revisit the abstraction principles behind the DUNE-FEM finite element library. Section 3 will be concerned with the restriction and prolongation of user data in *hp*- adaptive simulations. In Section 4 we describe the implementation and usage of the `dune-fem-hpdg` module. As proof of concept we implemented an *hp*-adaptive interior penalty method for an elliptic model problem, the numerical results are shown in Section 5.

2 Discrete functions and discrete function spaces

Let $\Omega \subset \mathbb{R}^d$ be a bounded domain. We assume that the computational domain is discretized by a grid $\mathcal{G} = \{E_i \subset \mathbb{R}^d \mid i \in I\}$, i.e., by a finite number of closed grid cells with non-overlapping interiors such that

$$\bigcup_{E \in \mathcal{G}} E = \overline{\Omega}.$$

A *discontinuous finite element space* or *discontinuous Galerkin space* is a finite-dimensional, piecewise continuous function space

$$X(\mathcal{G}) = \{u \in L^\infty(\Omega) \mid u\restriction_E \in X(E) \text{ for all } E \in \mathcal{G}\},$$

where for each $E \in \mathcal{G}$ we fixed a finite-dimensional subspace, the so-called local space $X(E) \subset C(E)$. For discretization purposes we need to fix a basis for each local space.

Definition 2.1 (Local basis function set)**.** *Let $X(E)$ be a local space with $\dim(X(E)) = n_E$. A basis*

$$\mathcal{B}_E = \{\varphi_{E,i} \mid i = 0, \dots, n_E - 1\}$$

of $X(E)$ is called a local basis function set*.*

A function $u \in X(\mathcal{G})$ is called a *discrete function*. Note that on inter-element intersections a discrete function is double-valued. The restriction of u to a grid element

$E \in \mathcal{G}$ is called a *local (discrete) function*. In the fixed local basis $\mathcal{B}_E$ a local function reads as follows,

$$u\restriction_E = \sum_{i=0}^{n_E-1} u_{E,i}\varphi_{E,i}.$$

The scalars $u_{E,i} \in \mathbb{R},\ i = 0, \dots, n_E - 1$, are called *local degrees of freedom* (DOFs); the vector $(u_{E,i})_{i=0,\dots,n_E-1}$ is called the *local DOF vector*. For practical purposes we want to store the vector of all DOFs in a consecutive array of floating point numbers. This requires a suitable enumeration of the DOFs.

Definition 2.2 (Local DOF mapping)**.** *Let $N = \dim(X(\mathcal{G}))$ denote the dimension of the finite element space $X(\mathcal{G})$. A* local DOF mapping *for $E \in \mathcal{G}$ is an injective mapping*

$$\mu_E : \{0, \dots, n_E - 1\} \to \{0, \dots, N - 1\}$$

from a local DOF enumeration to global indices.

Having fixed families of local basis function sets and local DOF mappings we have at the same time fixed a basis of $X(\mathcal{G})$.

Definition 2.3 (Global basis function set)**.** *Let $(\mathcal{B}_E)_{E\in\mathcal{G}}$ and $(\mu_E)_{E\in\mathcal{G}}$ be fixed families of local basis function sets and DOF mappings. Define $\psi_i \in X(\mathcal{G}),\ i = 1, \dots, N - 1$, by*

$$\psi_i\restriction_E = \begin{cases} \varphi_{E,j} & \text{if } \mu_E(j) = i, \\ 0 & \text{otherwise} \end{cases} \qquad (E \in \mathcal{G}).$$

Then, the set $\mathcal{B}_\mathcal{G} = \{\psi_i \mid i = 0, \dots, N - 1\}$ forms a basis of $X(\mathcal{G})$.

The following definition is a slightly modified version of [10, Definition 5] which represents the theoretical foundation for the implementations of finite element spaces spaces in Dune-Fem.

Definition 2.4 (Discrete function space)**.** *Let $X(\mathcal{G})$ be a finite-dimensional, piecewise continuous function space. Let further $(\mathcal{B}_E)_{E\in\mathcal{G}}$ and $(\mu_E)_{E\in\mathcal{G}}$ be families of local basis function sets and injective local DOF mappings,*

$$\mathcal{B}_E = \{\varphi_{E,i} \mid i = 0, \dots, n_E - 1\},$$
$$\mu_E : \{0, \dots, n_E - 1\} \to \{0, \dots, \dim(X(\mathcal{G})) - 1\},$$

such that $X(\mathcal{G}) = \operatorname{span}\mathcal{B}_\mathcal{G}$ with the global basis function set $\mathcal{B}_\mathcal{G}$ as defined in 2.3. Then, the triple

$$\left(X(\mathcal{G}), (\mathcal{B}_E)_{E\in\mathcal{G}}, (\mu_E)_{E\in\mathcal{G}}\right)$$

is called a discrete function space.

It is important to note that while we motivated the concept of a discrete function space starting from a discontinuous finite element space Definition 2.4 applies to continuous finite element spaces as well. For the sake of readability, we will not always explicitly state the local basis function sets and DOF mappings when referring to a discrete function space. Instead, we simply speak of a discrete function space $X(\mathcal{G})$ and it is tacitly understood that local basis function sets $(\mathcal{B}_E)_{E\in\mathcal{G}}$ and DOF mappings $(\mu_E)_{E\in\mathcal{G}}$ have been fixed.

Example 2.5. *Let $\boldsymbol{k} = (k_E)_{E\in\mathcal{G}}$ be a vector of local polynomial degrees. Denote by*

$$X^{\boldsymbol{k}}(\mathcal{G}) = \left\{u \in L^\infty(\Omega) \mid u\restriction_E \in P^{k_E}(E) \text{ for all } E \in \mathcal{G}\right\}$$

the space of piecewise polynomial functions. Assume that for each element $E \in \mathcal{G}$ there is a fixed reference element R_E and an affine reference mapping $F_E : R_E \to E$. We choose $\mathcal{B}_E = \mathcal{B}_{R_E,k_E} \circ F_E^{-1}$, where $\mathcal{B}_{R,k} \subset C(R)$ denotes the polynomial basis resulting from a Gram-Schmidt orthonormalization with respect to the $L^2(R)$ inner product applied to the lexicographically ordered monomials $(x^{\boldsymbol{\alpha}})_{|\boldsymbol{\alpha}|\leq k}$,

$$x^{\boldsymbol{\alpha}} = \prod_{i=1}^{d} x_i^{\alpha_i} \quad (\boldsymbol{\alpha} \in \mathbb{N}^d).$$

Note that for the vector space dimensions it holds

$$\dim(X^{\boldsymbol{k}}(\mathcal{G})) = \sum_{E\in\mathcal{G}} n_E = \sum_{E\in\mathcal{G}} \binom{k_E + d}{d}.$$

Assume there is a given strict total order $<$ on $\mathcal{G}$, e.g., from an enumeration of the grid elements. Then for each $E \in \mathcal{G}$ a local DOF mapping is provided by

$$\mu_E(i) = \sum_{\substack{E'\in\mathcal{G},\\ E'<E}} n_{E'} + i \quad (i = 0, \dots, n_E - 1).$$

3 Data transfer in *hp*-adaptive computations

Adaptive finite element methods give raise to sequences of approximate solutions, grids, and discrete function spaces. The main challenge we want to address in this section is the restriction and prolongation of user data in *hp*-adaptive simulations, i.e., any transfer of discrete functions and other grid-based data while modifying the mesh or the local ansatz spaces.

In the following, let $\left(X^{(m)}\right)_{m\in N}$ denote a sequence of discontinuous discrete function spaces,

$$X^{(m)} = \left\{u \in L^\infty(\Omega) \mid u\restriction_E \in X^{(m)}(E) \subset C(E) \text{ for all } E \in \mathcal{G}^{(m)}\right\}.$$

We assume that the associated sequence of grids $\left(\mathcal{G}^{(m)}\right)_{m\in\mathbb{N}}$ is nested in the following sense: for all $E^{(m+1)} \in \mathcal{G}^{(m+1)}\backslash\mathcal{G}^{(m)}$ we assume that either

i) there is a unique father element $E^{(m)} \in \mathcal{G}^{(m)}$ with $E^{(m+1)} \subset E^{(m)}$; in this case we say that $E \in \mathcal{G}^{(m+1)}$ resulted from refining $E^{(m)} \in \mathcal{G}(m)$,

ii) or there are a number of elements $E_j^{(m)} \in \mathcal{G}^{(m)}$, $j = 1, \dots, J$, such that $E^{(m+1)} = \bigcup_{j=1}^m E_j^{(m)}$; in this case we say that $E \in \mathcal{G}^{(m+1)}$ resulted from coarsening the children $E_j^{(m)}$, $j = 1, \dots, J$.

Readers familiar with DUNE and the definitions in [6] may think of the slightly more general case of sequences of codimension 0 leaf entity complexes in hierarchical meshes.

For each $m \in \mathbb{N}$ we fix a family of local projection operators

$$\Pi_E^{(m)} : L^\infty(E) \to X^{(m)}(E) \quad \left(E \in \mathcal{G}^{(m)}\right).$$

The most important example in view of discontinuous finite element methods is the following.

Example 3.1 (Local L^2-projection)**.** *For $m \in \mathbb{N}$ and $E \in \mathcal{G}^{(m)}$ the local $L^2(E)$-projection $\Pi_E^{(m)}$ is defined by*

$$\int_E (\Pi_E^{(m)} u)\varphi^{(m)} \,\mathrm{d}x = \int_E u\varphi^{(m)} \,\mathrm{d}x \quad \left(\varphi^{(m)} \in X^{(m)}(E)\right).$$

Having fixed a family of local projection operators we denote by $\Pi^{(m)} : L^\infty(\Omega) \to X^{(m)}$ the global projection operator defined by

$$\Pi^{(m)} u\restriction_E = \Pi_E^{(m)} u\restriction_E \quad \left(E \in \mathcal{G}^{(m)}\right)$$

for all $u \in X^{(m)}$. Now, let $u_0 \in L^\infty(\Omega)$ be some given initial data. On an abstract level, an adaptive scheme can be written in the following form,

$$\begin{aligned} u^{(0)} &= \Pi^{(0)} u_0, \\ u^{(m+1)} &= \Pi^{(m+1)} \circ \Phi^{(m)}(u^{(m)}) \quad (m \in \mathbb{N}), \end{aligned}$$

where each $\Phi^{(m)} : X^{(m)} \to X^{(m)}$ is some arbitrary operator. For the remainder of this section we will be concerned with the restriction and prolongation of a discrete function $u^{(m)} \in X^{(m)}$, i.e., the efficient computation of the projection

$$u^{(m+1)} = \Pi^{(m+1)} u^{(m)}.$$

From a mathematical point of view the restriction and prolongation of a given function is a trivial task. Its practical implementation, however, is not, the main difficulty being that during the modification data may be invalidated. In Dune, the grid adaptation is split into several stages. During this modification phase user data may be transferred from a grid state $\mathcal{G}^{(m)}$ to $\mathcal{G}^{(m+1)}$. For this particular purpose, a Dune grid provides persistent id mappings and associative containers, see [5]. Grid data stored as consecutive arrays (e.g., global DOF vectors) must be copied into temporary data structures that remain valid during the adaptation.

The finite element library Dune-Fem pursues a different strategy. During the modification phase DOF vectors are resized to hold information associated with elements $E \in \mathcal{G}^{(m)} \cup \mathcal{G}^{(m+1)}$. Each stage of the adaptation cycle requires updates on index sets, DOF mappings, and DOF vectors [10, Algorithms 18sqq.]. In case a relatively small number of elements is marked for local mesh adaptation this approach leads to significantly less memory overhead during adaptation. A comparison of the two different adaptation strategies can be found in [25]. Unfortunately, the algorithms and data structures implemented in Dune-Fem do not hold in case of *p*- or *hp*-adaptation. Their generalization must be done with great care if legacy code shall be supported. For the sake of presentation, we introduce the following notation.

Definition 3.2 (Global DOF set)**.** *For all $m \in \mathbb{N}$ we denote the dimension of the discrete function space $X^{(m)}$ by*

$$N^{(m)} = \dim\left(X^{(m)}\right) = \sum_{E \in \mathcal{G}^{(m)}} n_E^{(m)},$$

where $n_E^{(m)} = \dim X^{(m)}(E)$. *We define the* global DOF set $\mathcal{D}^{(m)} \subset \mathcal{G}^{(m)} \times \mathbb{N}$ *associated with* $X^{(m)}$ *by*

$$(E,i) \in \mathcal{D}^{(m)} :\Leftrightarrow i \in \{0,\dots,n_E^{(m)}-1\}.$$

A *global DOF mapping* is an injective mapping $\mu^{(m)} : \mathcal{D}^{(m)} \to \{0,\dots,N^{(m)}-1\}$. Obviously, a global DOF mapping $\mu^{(m)}$ is equivalent to a family of local DOF mappings $\left(\mu_E^{(m)}\right)_{E\in\mathcal{G}^{(m)}}$ by the relation

$$\mu_E^{(m)}(i) = \mu^{(m)}(E,i) \quad \left((E,i) \in \mathcal{D}^{(m)}\right). \tag{1}$$

The following definition describes our generalized approach to the restriction and prolongation of discrete functions.

Definition 3.3 (Restriction and prolongation)**.** *Assume a global DOF mapping* $\mu^{(m)}$ *is already known. Let*

$$u^{(m)} = \sum_{i=0}^{N^{(m)}-1} u_i \psi_i^{(m)} \in X^{(m)}$$

be a given discrete function developed in the global basis $\mathcal{B}_\mathcal{G}$ *defined in 2.3. The global DOF vector will be denoted by* $\boldsymbol{u} = (u_i)_{i=0,\dots,N^{(m)}-1}$, *and we deliberately omit the index* m. *In order to simultaneously compute a global DOF mapping* $\mu^{(m+1)}$ *and the projection* $u^{(m+1)} = \Pi^{(m+1)}u^{(m)}$ *proceed as follows.*

Step 1 (Insertion of new DOFs). *Let* $\mathcal{D}^{(m+1/2)} = \mathcal{D}^{(m)} \cup \mathcal{D}^{(m+1)}$ *and* $\mu^{(m+1/2)}$ *be a continuation of* $\mu^{(m)}$ *on* $\mathcal{D}^{(m+1/2)}$, *i.e.,* $\mu^{(m+1/2)}$ *is an injective mapping*

$$\mu^{(m+1/2)} : \mathcal{D}^{(m+1/2)} \to \{0,\dots,|\mathcal{D}^{(m+1/2)}|-1\},$$

such that

$$\mu^{(m+1/2)}(E,i) = \mu^{(m)}(E,i) \quad \left((E,i) \in \mathcal{D}^{(m)}\right).$$

Step 2 (Restriction and prolongation of user data). *The global DOF vector* $\boldsymbol{u}$ *is temporarily resized to* $N^{(m+1/2)} = |\mathcal{D}^{(m+1/2)}|$. *In case of a newly inserted element* $E \in \mathcal{G}^{(m+1)}\backslash\mathcal{G}^{(m)}$ *created either from local grid refinement or coarsening the associated DOFs are initialized by*

$$\sum_{i=0}^{n_E^{(m+1)}-1} u_{\mu^{(m+1/2)}(E,i)} \varphi_{E,i}^{(m+1)} = \Pi^{(m+1)} u^{(m)}\restriction_E.$$

If otherwise $X^{(m+1)}(E) \neq X^{(m)}(E)$ *all DOFs* $u_{\mu^{(m+1/2)}(E,i)}$, $i = 0,\dots,n_E^{(m+1)}-1$, *must be reinitialized by projecting the local discrete function* $u^{(m)}\restriction_E$ *to the local space* $X^{(m+1)}(E)$. *Note that in general* $\mathcal{D}_E^{(m)} \cap \mathcal{D}_E^{(m+1/2)} \neq \emptyset$. *We make a copy of the local DOF vector,*

$$v_i = u_{\mu^{(m+1/2)}(E,i)} \quad (i = 0,\dots,n_E^{(m)}-1),$$

and compute the projection

$$\sum_{i=0}^{n_E^{(m+1)}-1} u_{\mu^{(m+1/2)}(E,i)} \varphi_{E,i}^{(m+1)} = \Pi_E^{(m+1)} \sum_{i=0}^{n_E^{(m)}-1} v_i \varphi_{E,i}^{(m)}.$$

Step 3 (Removal of DOFs). *Construct a new injective DOF numbering*

$$\mu^{(m+1)} : \mathcal{D}^{(m+1)} \to \{0, \dots, N^{(m+1)} - 1\},$$

such that

$$\mu^{(m+1)}(E, i) = \mu^{(m+1/2)}(E, i) \quad \text{if } \mu^{(m+1/2)}(E, i) < N^{(m+1)}.$$

All other DOFs are copied to their new destination,

$$u_{\mu^{(m+1)}(E,i)} = u_{\mu^{(m+1/2)}(E,i)} \quad \text{if } \mu^{(m+1/2)}(E, i) \geq N^{(m+1)}$$

and the DOF vector $\boldsymbol{u}$ *is resized to its new length* $N^{(m+1)}$.

The algorithm 3.3 yields a global DOF mapping $\mu^{(m+1)}$ and by Equation (1) the family of associated local DOF mappings $\left(\mu_E^{(m+1)}\right)_{E \in \mathcal{G}^{(m+1)}}$. At the same time we have computed the projection of $u^{(m)}$ onto $X^{(m+1)}$,

$$u^{(m+1)} = \sum_{i=0}^{N^{(m+1)}-1} u_i \psi_i^{(m+1)},$$

developed in the global basis functions defined in 2.3.

Concerning the actual implementation of the above algorithm, Steps 1 and 2 should be clear. The practical definition of the DOF mapping $\mu^{(m+1)}$ in Step 3, however, may be in need of further explanation.

Example 3.4 (Removal of DOFs)**.** *Let* $\mu^{(m+1/2)}$ *be the intermediate DOF mapping as in Step 1 above. We denote by*

$$\mathcal{H}^{(m+1/2)} = \{0, \dots, N^{(m+1)} - 1\} \backslash \mu^{(m+1/2)}(\mathcal{D}^{(m+1)})$$

the set of freed valid indices, the so-called set of holes. Its elements h_i, $i = 0, \dots, |\mathcal{H}^{(m+1/2)}| - 1$, *are assumed to be in ascending order. Let*

$$\mu^{(m+1)} : \mathcal{D}^{(m+1)} \to \{0, \dots, |\mathcal{D}^{(m+1)}| - 1\}$$

be defined by

$$\mu^{(m+1)}(E, j) = \begin{cases} \mu^{(m+1/2)}(E, j) & \text{if } \mu^{(m+1/2)}(E, j) < |\mathcal{D}^{(m+1)}| - 1, \\ h_i & \text{if } \mu^{(m+1/2)}(E, j) = |\mathcal{D}^{(m+1)}| + k_i, \end{cases}$$

where the numbers k_i, $i = 0, \dots, |\mathcal{H}^{(m+1/2)}| - 1$, *are assumed to be in ascending order as well.*

Example 3.5. *Figure 1 shows a hierarchical grid with three macro elements* E_0, E_1, E_2 *with* E_0 *having two children* E_3, E_4. *The leaf level elements form the grid* $\mathcal{G}^{(m)} = \{E_i \mid i = 1, \dots, 4\}$; *the global DOF numbers associated with each element are depicted inside the brackets. Here, level* 0 *elements are assumed to hold at least two DOFs, while level* 1 *elements shall have one DOF. Next,* E_3 *and* E_4 *are marked for coarsening,* E_2 *is marked for refinement, and* E_1 *gets assigned an additional DOF by local p-refinement. Newly inserted DOFs added in Step 1 and DOFs to be removed in Step 3 are printed in italics. During the modification phase the number of DOFs is enlarged to allow for the local projection of user data.*

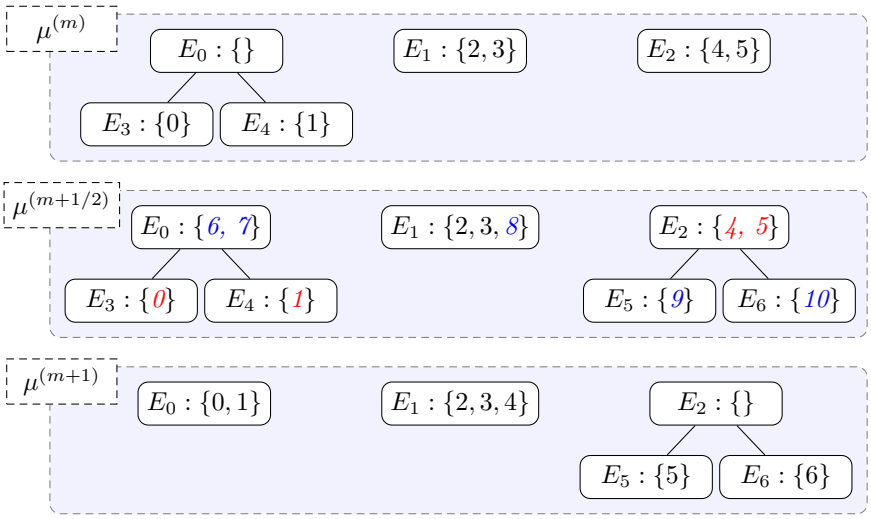

Figure 1: Computation of a global DOF mapping $\mu^{(m+1)}$ by Steps 1 to 3 of Definition 3.3, see Example 3.5 for details.

Remark 3.6 (Storage costs)**.** *During the adaptation from $X^{(m)}$ to $X^{(m+1)}$ the maximum length of the global DOF vector equals $|\mathcal{D}^{(m+1/2)}| = \mathcal{O}(|\mathcal{G}^{(m)} \cup \mathcal{G}^{(m+1)}|)$. Taking into account the temporary storage used in Step 2 the overall memory consumption of algorithm 3.3 is of order $\mathcal{O}(|\mathcal{G}^{(m)} \cup \mathcal{G}^{(m+1)}| + 1)$.*

4 Using the dune-fem-hpdg module

In this section we describe the usage of the `dune-fem-hpdg` add-on module to the DUNE-FEM finite element library. The module provides extensible reference implementations of adaptive discrete function spaces for implementing *p*- and *hp*-adaptive discontinuous finite element methods. First, we give a brief introduction to DUNE-FEM and in particular its local mesh adaptation capabilities.

4.1 The Dune-Fem finite element library

The C++-library DUNE-FEM provides a number of abstract interface classes representing discrete functions and discrete function spaces, local basis function sets, and DOF vectors and DOF mappings. These classes and their relationships are summarized in Figure 2. Here, a `DiscreteFunction` represents an element $u \in X(\mathcal{G})$. A discrete function holds a global DOF vector

$$\boldsymbol{u} = (u_0, \dots, u_{N-1})$$

and is usually evaluated in local coordinates with respect to a given element $E \in \mathcal{G}$. The local DOF vector

$$(u_{\mu_E(0)}, \dots, u_{\mu_E(n_E-1)})$$

is initialized by the `DofMapper`, a class representing the family of local DOF mappings $(\mu_E)_{E\in\mathcal{G}}$. Finally, a `BasisFunctionSet` represents the local basis function set $\mathcal{B}_E$, and

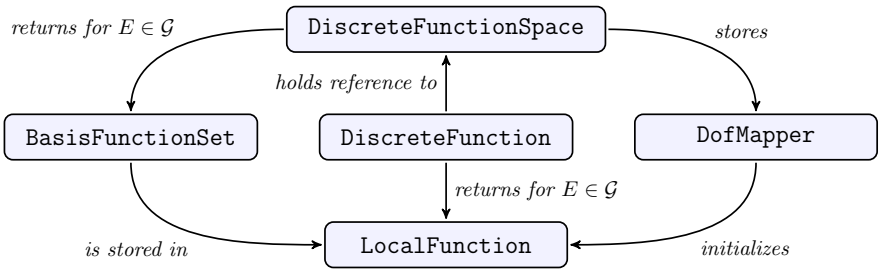

Figure 2: Interplay of the main interface classes in DUNE-FEM representing discrete function spaces and discrete functions.

Listing 1: Sample code illustrating the automated restriction and prolongation of discrete functions under grid refinement in DUNE-FEM.

```
template <class DiscreteFunction>
void adapt(DiscreteFunction &uh, unsigned int seed) {
  // get grid (slightly simplified)
  using Grid = typename DiscreteFunction::GridType;
  Grid &grid = uh.grid();

  // create adaptation manager (slightly simplified)
  using RestrictProlong = RestrictProlongDefault<DiscreteFunction>;
  using h_AdaptationManager = AdaptationManager<Grid, RestrictProlong>;
  h_AdaptationManager h_adaptManager(RestrictProlong(uh));

  // randomly mark leaf level grid elements for refinement/coarsening
  default_random_engine engine(seed);
  uniform_int_distribution<int> distribution(-1, 1);
  for (const auto &element : elements(grid.leafGridView()))
    grid.mark(distribution(engine), element);

  // adapt discrete function
  h_adaptManager.adapt();
}
```

the local function $u\restriction_E \in X(E)$ is given by

$$u\restriction_E = \sum_{i=0}^{n_E-1} u_{\mu_E(i)} \varphi_{E,i}.$$

In most complex applications, more than one discrete function is in use. This poses some difficulty in adaptive simulations. A grid may be only be adapted once, and all discrete functions must be restricted and prolonged simultaneously. In DUNE-FEM, a sole class, the `DofManager`, is responsible for starting the grid adaptation process. It resizes all DOF vectors and notifies DOF mappings about the update.

In case no user data must be restricted and prolonged no action is required. Otherwise, users must add discrete functions, containers, etc., to an `AdaptationManager`. This class is equipped with the necessary information (encapsulated in a `RestrictProlong` object) on how to locally project user data during adaptation. Listing 1 illustrates the adaptation process from a user perspective.

4.2 Discrete function spaces in dune-fem-hpdg

As illustrated in Figure 2 a `DiscreteFunctionSpace` assembles the local basis function sets and DOF mappings. The `dune-fem-hpdg` extends the list of available discrete function spaces available to users of DUNE-FEM by the following *hp*-adaptive spaces:

`OrthogonalDiscontinuousGalerkinSpace` an implementation of the example finite element space 2.5,

`LegendreDiscontinuousGalerkinSpace` a discontinuous finite element space using product Legendre ansatz polynomials, i.e., $\mathcal{B}_{E,k} = \{\varphi_{\boldsymbol{\alpha}} \mid 0 \leq \alpha_i \leq k, i = 1, \dots, d\}, k \in \mathbb{N}$, with

$$\varphi_{\boldsymbol{\alpha}}(F_E(x)) = \prod_{i=1}^{d} p_{\alpha_i}(x_i) \quad (x \in R_E),$$

where $F_E : R_E \to E$ is an affine reference mapping and the i-th Legendre polynomial on the unit interval $[0, 1]$ is defined by

$$p_i(x) = \frac{1}{i!} \frac{\mathrm{d}^i}{\mathrm{d}x^i} \left[(x^2 - x)^i\right],$$

`AnisotropicDiscontinuousGalerkinSpace` a discontinuous space based on product Legendre polynomials as well; this implementation, however, allows for an adaptation in p in each spatial direction.

We remark that the latter two spaces based on Legendre polynomials are restricted to cubic grids, i.e., $R_E = [0, 1]^d$ for all $E \in \mathcal{G}$.

The discrete function spaces listed above differ in the choice of local basis function sets. They do, however, share the common base class `hpDG::DiscontinuousGalerkinSpace` providing almost all other functionality. In the same way, the `dune-fem-hpdg` module allows users to quickly setup new discrete function spaces. In order to do so, users must provide an implementation of the `hpDG::BasisFunctionSets` interface which represents a family of local basis function sets. This is explained in more detail in Appendix B.

The construction of the discontinuous finite element spaces in `dune-fem-hpdg` is in no way different from that of all other discrete function spaces in DUNE-FEM. Listing 2 illustrates the initialization of an adaptive discrete function suitable for *hp*-adaptive simulations.

4.3 An interface for local *p*-adaptation

The most important design decision we made in implementing the `dune-fem-hpdg` module was to split *h*- and *p*-refinement in two separate stages. Our reasoning is twofold. First, the software should support *h*-adaptive legacy code in the manner described above in Section 4.1. Second, this strategy allows for an arbitrary number of *h*-, *p*- and *hp*-adaptive discrete functions in single application and gives users the necessary freedom in complex applications. Standard use cases are easy to implement as will be shown in Section 5.

Any of the discrete function spaces in `dune-fem-hpdg` provides a set of extended interface methods for local *p*-adaptation. In most papers, p denotes an integer, e.g., the polynomial order of a local approximation space. We slightly generalized this idea and allow p to be of arbitrary type (the `Key` type), e.g., a vector of individual local polynomial degrees for each space direction. The *p*-adaptive interface mimicks that of the grid adaptation in DUNE.

Listing 2: Construction of a discrete function space and a discrete function with `dune-fem-hpdg` suitable for *hp*-adaptive simulations.

```
// create quadrilateral ALUGrid
using Grid = ALUGrid<2, cube, nonconforming>;
Grid grid(/*...*/);

// select leaf level partition
using GridPart = AdaptiveLeafGridPart<Grid>;
GridPart gridPart(grid);

// create scalar discrete function space
const int order = 3;
using DiscreteFunctionSpace = hpDG::OrthogonalDiscontinuousGalerkinSpace<
    FunctionSpace<double, double 2, 1>, GridPart, order>;
DiscreteFunctionSpace space(gridPart);

// create discrete function
using DiscreteFunction = AdaptiveDiscreteFunction<DiscreteFunctionSpace>;
DiscreteFunction uh("uh", space);
```

Listing 3: Sample code illustrating the automated restriction and prolongation of discrete functions under modification of the local polynomial degree.

```
template <class DiscreteFunction>
void adapt(DiscreteFunction &uh, unsigned int seed) {
  // get space (slightly simplified)
  using DiscreteFunctionSpace =
      typename DiscreteFunction::DiscreteFunctionSpaceType;
  DiscreteFunctionSpace &space = uh.space();

  // create adaptation manager
  using Grid = typename DiscreteFunction::GridType;
  using DataProjection = hpDG::DefaultDataProjection<DiscreteFunction>;
  using p_AdaptationManager = hpDG::AdaptationManager<Grid,
      DataProjection>;
  p_AdaptationManager p_adaptManager(DataProjection(uh));

  // randomly mark grid (part) elements for p-adaptation
  default_random_engine engine(seed);
  uniform_int_distribution<int> distribution(0, space.order());
  for (const auto &element : elements(uh.gridPart()))
    space.mark(distribution(engine), element);

  // adapt discrete function
  p_adaptManager.adapt();
}
```

```
const DiscreteFunctionSpace::KeyType &key(
    const DiscreteFunctionType::EntityType &entity) const;
void DiscreteFunctionSpace::mark(
    const DiscreteFunctionType::KeyType &key,
    const DiscreteFunctionType::EntityType &element);
```

The first method returns the key currently assigned to a grid element, while `mark` allows for re-assigning a new key. In order for the changes to take effect, one of the following two `adapt` methods must be called.

```
bool DiscreteFunctionSpace::adapt();
template <DataProjection>
bool DiscreteFunctionSpace::adapt(DataProjection &projection);
```

The necessary information on how to locally project data are encapsulated in a `DataProjection`. As in case of h-adaptation, however, users will not call these methods explicitly. Instead, the `hpDG::AdaptationManager` class handles the restriction and prolongation as illustrated by Listing 3. Note how closely the code resembles Listing 1 above on local mesh adaptation in DUNE-FEM. In this particular example, a single discrete function is restricted and prolonged; however, the `dune-fem-hpdg` facilities may be provided with an arbitrary number of discrete functions and other user data. More information on the adaptation and the restriction and prolongation of custom user data can be found in Appendix B.

5 An *hp*-adaptive interior penalty Galerkin method

In this final section we want to illustrate in a complex application the capabilities and usage of the `dune-fem-hpdg` module. We will be concerned with the numerical solution of an elliptic PDE by means of an *hp*-adaptive symmetric interior penalty Galerkin (SIPG) scheme. First, we briefly revisit the *hp*-version of the SIPG method following the recent book by Dolejší and Feistauer [13, Chapter 7].

5.1 The *hp*-version of the SIPG method

Let $\Omega \subset \mathbb{R}^2$ be a bounded domain with Lipschitz boundary $\partial\Omega$. We consider the following elliptic model problem

$$-\Delta u = f \quad \text{in } \Omega, u \quad = g \quad \text{on } \partial\Omega \tag{2}$$

for given source f and boundary values g. The domain is discretized by a grid $\mathcal{G}$; for $E \in \mathcal{G}$ we denote the sets of interior and boundary intersections by

$$\mathcal{I}_{int}(E) = \{e = E \cap E' \mid E' \in \mathcal{G}\backslash\{E\}\},$$
$$\mathcal{I}_{bnd}(E) = \{e = \partial E \cap \partial\Omega\}.$$

The sets of all interior and boundary intersections will be denoted by Γ_{int} and Γ_{bnd}, respectively. We fix local polynomial degrees $\boldsymbol{k} = (k_E)_{E\in\mathcal{G}}$ with $k_E \geq 1$ for all $E \in \mathcal{G}$ and consider the standard discontinuous finite element space

$$X^{\boldsymbol{k}}(\mathcal{G}) = \{u \in L^\infty(\Omega) \mid u\restriction_E \in P^{k_E}(E) \text{ for all } E \in \mathcal{G}\}.$$

The jump of a discrete function $u \in X^{\boldsymbol{k}}(\mathcal{G})$ across an inter-element intersection $e \in \Gamma_{int}$, $e = E \cap E'$, is defined by

$$[\![u]\!](x) = \left(u\restriction_E \nu_E + u\restriction_{E'} \nu_{E'}\right)(x) \quad (x \in e),$$

where ν_E denotes a unit outer normal to $E \in \mathcal{G}$. The average of a discrete function is given by

$$\{u\}(x) = \frac{1}{2}\big(u\restriction_E + u\restriction_{E'}\big)(x) \quad (x \in e).$$

For each intersection $e \in \Gamma_{int} \cup \Gamma_{bnd}$ we introduce the penalty parameter

$$\sigma_e = \begin{cases} \gamma \dfrac{k_E^2 + k_{E'}^2}{2h_e} & \text{for } e = E \cap E', \\ \gamma \dfrac{k_E^2}{h_e} & \text{for } e = \partial E \cap \partial\Omega, \end{cases}$$

where $\gamma > 0$ is a sufficiently large constant and h_e denotes the diameter of e. Note that the penalty parameter depends on the local polynomial degrees. Now, let $B : X^{\boldsymbol{k}}(\mathcal{G}) \times X^{\boldsymbol{k}}(\mathcal{G}) \to \mathbb{R}$ be a bilinear form defined by

$$\begin{aligned} B(u,\psi) = &\sum_{E\in\mathcal{G}} \int_E \nabla u \cdot \nabla\psi \,\mathrm{d}x \\ &- \sum_{e\in\Gamma_{int}} \int_e [\![u]\!]\{\nabla\psi\} + [\![\psi]\!]\{\nabla u\} \,\mathrm{d}s + \sum_{e\in\Gamma_{int}} \int_e \sigma_e [\![u]\!][\![\psi]\!] \,\mathrm{d}s \\ &- \sum_{e\in\Gamma_{bnd}} \int_e u(\nabla\psi\cdot\nu) + \psi(\nabla u\cdot\nu) \,\mathrm{d}s + \sum_{e\in\Gamma_{bnd}} \int_e \sigma_e u\psi \,\mathrm{d}s, \end{aligned}$$

and $l : X^{\boldsymbol{k}}(\mathcal{G}) \to \mathbb{R}$ defined by

$$l(\psi) = \int_\Omega f\psi \,\mathrm{d}x - \sum_{e\in\Gamma_{bnd}} \int_e g(\nabla\psi\cdot\nu) \,\mathrm{d}s + \sum_{e\in\Gamma_{bnd}} \int_e \sigma_e g\psi \,\mathrm{d}s.$$

The bilinear form B is continuous and coercive with respect to the energy norm $\|\cdot\|_{DG} : X^{\boldsymbol{k}}(\mathcal{G}) \to \mathbb{R}$ defined by

$$\|u\|_{DG}^2 = \sum_{E\in\mathcal{G}} \int_E |\nabla u|^2 \,\mathrm{d}x + \sum_{e\in\Gamma_{int}} \int_e \sigma_e [\![u]\!]^2 \,\mathrm{d}s + \sum_{e\in\Gamma_{bnd}} \int_e \sigma_e u^2 \,\mathrm{d}s$$

provided the constant γ is sufficiently large [13, Theorems 7.13 and 7.15]. This guarantees the existence and uniqueness of the weak solution $u_h \in X^{\boldsymbol{k}}(\mathcal{G})$ to the model problem (2) defined by

$$B(u_h,\psi) = l(\psi) \quad (\psi \in X^{\boldsymbol{k}}(\mathcal{G})). \tag{3}$$

5.2 The *hp*-adaptive scheme

The first component in the construction of a fully *hp*-adaptive scheme is an *a posteriori* error indicator estimating the local approximation error. Of course, the local error indicators must be computable from the numerical solution u_h and the given data f, g alone. There is a growing body of literature on *a posteriori* error estimation, see, e.g.,

[12, 21, 32]. We implemented the following indicator from [24],

$$\begin{aligned}\eta_E^2 =& \frac{{h_E}^2}{{k_E}^2}\int_E \big(\Pi_{E,k_E-1}(f+\Delta u_h)\big)^2\,\mathrm{d}x \\ &+ \sum_{e\in\mathcal{I}_{int}(E)} \left\{ \frac{h_E}{k_E}\int_e \big(\Pi_{e,k_e-1}[\![\nabla u_h\cdot\nu_E]\!]\big)^2\,\mathrm{d}s + \frac{{k_E}^3}{h_E}\int_e [\![u_h]\!]^2\,\mathrm{d}s \right\} \\ &+ \sum_{e\in\mathcal{I}_{bnd}(E)} \frac{{k_E}^3}{h_E}\int_e (u_h-g)^2\,\mathrm{d}s \quad (E\in\mathcal{G}).\end{aligned} \tag{4}$$

Here, Π_{E,k_E-1} and Π_{e,k_e-1} denote local L^2-projections onto the polynomial spaces of lower degree $P^{k_E-1}(E)$ and $P^{k_e-1}(e)$ with $k_e=\max\{k_E,k_{E'}\}$. It can be shown that

$$\|u-u_h\|_{DG} \le \left(\sum_{E\in\mathcal{G}}\eta_E^2 + \mathcal{O}(f,u_h)\right)^{\frac{1}{2}},$$

where $\mathcal{O}(f,u_h)$ denotes a data-oscillation term, see [24, Theorem 3.2] for details. Having computed the local error indicator we mark $E\in\mathcal{G}$ for refinement in either h or p if the local error indicator exceeds an upper bound η^*,

$$\eta_E < \eta_* = TOL/|\mathcal{G}|.$$

If no element is marked for further refinement we stop the iterative procedure.

Once a grid element has been identified for local adaptation it must be decided whether to refine in h or p. Several strategies for this have been proposed in the literature, see, e.g., [11, 14, 22, 23, 29, 31]. We implemented the so-called PRIOR2P strategy described in [28] based on the following idea. *A priori* error estimates (see, e.g., [13, Theorem 7.20]) suggest that we should increase the local polynomial degree k_E provided the exact is sufficiently smooth in $E\in\mathcal{G}$. A regularity indicator yields an estimate for the local Sobolev index q_E,

$$q_E = \max\{q \mid u\restriction_E \in H^q(E)\} \quad (E\in\mathcal{G}).$$

We increase the local polynomial order provided $k_E < q_E - 1$; otherwise, we mark E for local mesh refinement.

So far we have only discussed local *hp*-refinement. Assume now that for $E\in\mathcal{G}$ the estimated local approximation error is small, i.e.,

$$\eta_E \le \eta_* \ll \eta^*.$$

We may decide to reduce the number of degrees of freedom by decreasing the local polynomial degree k_E or by coarsening the grid. Mesh coarsening usually depends on a number of neighboring elements (e.g., the children of a common father element in a hierarchical grid) all being marked for grid coarsening. We take a rather hands-on approach and — if possible — always favor h- over p-coarsening. We remark that the regularity indicator implemented restricts the minimum local polynomial degree to $k_{min}=3$.

Finally, we need to address the restriction and prolongation of the vector of local polynomial degrees $\boldsymbol{k}$. Let $E\in\mathcal{G}^{(m+1)}\backslash\mathcal{G}^{(m)}$ be an element created during grid modification. We must define the local polynomial degree $k_E^{(m+1)}$ to be associated with the

Algorithm 5.1 Summary of the *hp*-adaptive SIPG method.

choose initial grid $\mathcal{G}$ and local polynomial degrees $\boldsymbol{k} = (k_E)_{E\in\mathcal{G}}$
for $m = 0, 1, \ldots$ **do**
  solve system (3) to compute $u_h \in X^{\boldsymbol{k}}(\mathcal{G})$
  for all elements $E \in \mathcal{G}$ **do**
    compute the local error indicator η_E from (4)
  end for
  if $\left(\sum_{E\in\mathcal{G}} \eta_E^2\right)^{\frac{1}{2}} \leq TOL$ **then**
    stop
  end if
  for all elements $E \in \mathcal{G}$ **do**
    if $\eta_E < \eta_*$ **then**
      mark element E for h-coarsening, if possible;
      otherwise, set $k_E \leftarrow \max\{k_E - 1, k_{min}\}$
    else if $\eta_E > \eta^*$ **then**
      compute estimate for the local Sobolev index q_E
      if $q_E > k_E + 1$ **then**
        set $k_E \leftarrow \min\{k_E + 1, k_{max}\}$
      else
        mark element E for h-refinement
      end if
    end if
  end for
  while adapting the grid $\mathcal{G}$ **do**
    restrict and prolong u_h and $\boldsymbol{k}$ according to (5) and (6)
  end while
end for

newly created element. In case E results from local mesh refinement of $E' \in \mathcal{G}^{(m)}$ we simply set

$$k_E^{(m+1)} = k_{E'}^{(m)} \tag{5}$$

locally prolonging $\boldsymbol{k}^{(m)}$ to $\boldsymbol{k}^{(m+1)}$. If otherwise E results from local mesh coarsening we locally restrict $\boldsymbol{k}^{(m)}$ to $\boldsymbol{k}^{(m+1)}$ by setting

$$k_E^{(m+1)} = \max_{\substack{E'\in\mathcal{G}^{(m)},\\ E'\subset E}} k_{E'}^{(m)}. \tag{6}$$

The complete *hp*-adaptive SIPG scheme is summarized in Algorithm 5.1.

5.3 Implementation details

The `dune-fem-hpdg` contains a reference implementation of the *hp*-adaptive SIPG scheme described above. For the computation of the numerical solution we relied on DUNE and the DUNE-FEM discretization module. The latter provides sample implementations of continuous and discontinuous finite element methods including an SIPG method for fixed global polynomial degree k. For the solution of the linear system (3) we used the Iterative Solver Template Library (DUNE-ISTL), a DUNE module developed by Blatt and Bastian [7, 8]. The *hp*-adaptive scheme requires for a locally adaptive grid. In all

Listing 4: Implementation of Algorithm 5.1, Lines 10 to 22. Note that the given vector of polynomial degrees k associated with the discontinuous finite element space $X^{(m)}$ is overridden.

```
enum class Flag { Coarsen, None, Refine };

template <class Grid, class ManagedArray, class Function>
void mark(Grid &grid, ManagedArray &k, Function function) {
  for (const auto &element : elements(grid.leafGridView())) {
    // get refinement flag and estimated Sobolev index
    Flag what;
    int q;
    tie(what, q) = function(element);

    // mark grid/vector of local polynomial degrees for
    // hp-refinement/coarsening
    switch (what) {
      case Flag::None:
        break;
      case Flag::Coarsen:
        if (element.hasFather())
          grid.mark(-1, element);
        else if (k[element] > 3)
          k[element] -= 1;
        break;
      case Flag::Refine:
        if (k[element] < q - 1)
          k[element] += 1;
        else
          grid.mark(1, element);
    }
  }
}
```

our experiments we used the DUNE-ALUGRID module by Alkämper et al. [1], a DUNE add-on encapsulating the ALUGrid library by Schupp [30].

We want to give some details on the implementation in order to illustrate the *hp*-adaptation process. Listing 4 shows a sample implementation of the marking procedure (Lines 10 to 22 of Algorithm 5.1).

Here, the parameter `function` is a callable object which returns a pair of a refinement flag and an estimate for the local Sobolev index for a given entity. The array `k` is an associative container of local polynomial degrees currently in use. In case an element is flagged for refinement or coarsening it is either marked for h- or p-adaptation; in the latter case we re-assign its local polynomial degree in `k`. The modification of the grid and the finite element space as well as the restriction and prolongation of the discrete data are shown in Listing 5.

The first input argument of `adapt` is the approximate solution $u_h^{(m)} \in X^{(m)} = X^{\boldsymbol{k}^{(m)}}(\mathcal{G}^{(m)})$. First, the grid is adapted to its new state $\mathcal{G}^{(m+1)}$. The approximate solution and the container are restricted $\boldsymbol{k}^{(m)}$ by the usual DUNE-FEM facilities. This yields the vector $\boldsymbol{k}^{(m+1)}$ and an intermediate projection $u_h^{(m+1/2)} \in X^{\boldsymbol{k}^{(m)}}(\mathcal{G}^{(m+1)})$. In Lines 20 and 21 the finite element space is marked for the concluding p-adaptation. The subsequent restriction and prolongation is handled by `dune-fem-hpdg`. It yields the desired

Listing 5: Prolongation of the vector of polynomial degrees during grid adaptation. It is assumed that the function `mark` from Listing 4 has been called in advance.

```
template <class DiscreteFunction, class ManagedArray>
void adapt(DiscreteFunction &u, ManagedArray &k) {
  // get discrete function space (slightly simplified)
  using DiscreteFunctionSpace =
      typename DiscreteFunction::DiscreteFunctionSpaceType;
  DiscreteFunctionSpace &space = u.space();

  // get hierarchical grid
  using Grid = typename DiscreteFunctionSpace::GridType;
  Grid &grid = space.grid();

  // h-adaptation, restrict/prolong numerical solution and polynomial
      degrees
  using RestrictProlong =
      RestrictProlongDefaultTuple<DiscreteFunction, ManagedArray>;
  using h_AdaptationManager = AdaptationManager<Grid, RestrictProlong>;
  h_AdaptationManager h_AdaptManager(grid, RestrictProlong(u, k));
  h_AdaptManager.adapt();

  // equip discrete function space with new local polynomial degrees
  for (const auto &element : elements(grid.leafGridView()))
    space.mark(k[element], element);

  // p-adaptation, restrict/prolong numerical solution
  using DataProjection = hpDG::DefaultDataProjection<DiscreteFunction>;
  using p_AdaptationManager =
      hpDG::AdaptationManager<DiscreteFunctionSpace, DataProjection>;
  p_AdaptationManager p_AdaptManager(space, DataProjection(u));
  p_AdaptManager.adapt();
}
```

projection of $u_h^{(m)}$ onto $X^{(m+1)}$. The resulting function serves as initial guess in the solution of the next linear system (3).

5.4 Numerical results

We consider the homogeneous model problem to recover a prescribed solution given by

$$u(x, y) = r^{2/3} \sin(2\varphi/3)$$

in a domain with reentrant corner $\Omega = (-1, 1)^2 \backslash (0, -1) \times (1, 0)$. We compute two series of numerical solutions $u_h^{(m)}$, $m = 0, \ldots, 8$, on an axis-aligned quadrilateral mesh and a triangular mesh. In each case the *hp*-adaptive SIPG scheme yields a sequence of locally adapted meshes and distributions of local polynomial degrees. The so-called *hp*-meshes are shown in Figures 4,5 and 6,7. Since the exact solution is known we can compute the approximation error

$$e_h^{(m)} = u - u_h^{(m)}.$$

In order to determine the convergence rate (EOC) we follow [13] and define

$$\text{EOC} = -\frac{\log(\|e_h^{(m)}\| / \|e_h^{(m+1)}\|)}{\log((N^{(m)}/N^{(m+1)})^{1/d})},$$

where $N^{(m)}$ denotes the number of global DOFs. Another interesting quantity for evaluating the effectiveness of the adaptive scheme is the so-called effectivity index, defined as the ratio of the *a posteriori* error indicator error bound and the energy norm, i.e.,

$$\text{Eff. index} = \frac{\left(\sum_{E \in \mathcal{G}^{(m)}} \eta_E^{(m)^2}\right)^{\frac{1}{2}}}{\|e_h^{(m)}\|_{DG}}.$$

The effectivity index depends on the problem under consideration and the macro grid. Within the elements of given series of approximate solutions, however, the effectivity index should be a constant. The approximation errors, convergence rates, and effectivity indices for the model problem are shown in Tables 1 and 2. We encourage users to run the sample code themselves; the instructions can be found in Appendix A.

A Installing and running the software

In this appendix we describe how to download and install the `dune-fem-hpdg` module and how to reproduce the numerical results presented in Section 5.

A.1 Download of required Dune modules

The following DUNE modules are needed in order to build and run the examples:

i) the DUNE 2.4 core modules DUNE-COMMON, DUNE-GEOMETRY, DUNE-GRID, DUNE-ISTL, and DUNE-LOCALFUNCTIONS[1],

ii) the DUNE-ALUGRID library[2],

[1] `http://www.dune-project.org/download.html`

[2] `https://gitlab.dune-project.org/extensions/dune-alugrid`

iii) the DUNE-FEM discretization module[3],

iv) and the `dune-fem-hpdg` module[4].

Please make sure that you check out the DUNE 2.4 compatible (release) branches of DUNE-ALUGRID, DUNE-FEM, and `dune-fem-hpdg`.

A.2 Installation of the software

Before installing DUNE and its components please refer to the installation notes[5] for a list of dependencies and required software. We assume all aforementioned DUNE packages have been downloaded and saved to a single directory `$DUNE`. For the simultaneous build of all modules using CMake run the `dunecontrol` script, e.g., by

```
cd $DUNE
./dune-common/bin/dunecontrol all
```

Please refer to the installation notes for more information and on how to pass options to the build process.

A.3 Running the *hp*-adaptive sample code

We assume that DUNE has been configured using CMake. By default, binary executables will be built out-of-source, e.g., in a designated build directory `build_cmake`. The sources for the *hp*-adaptive SIPG method can be found in the subdirectory `examples/poisson`. In order to verify the results shown in Figures 4, 5 and Table 1 compile and run the test as follows

```
cd $DUNE/dune-fem-hpdg/build-cmake
cd ./examples/poisson
make poisson_alugrid_cube_7
./poisson_alugrid_cube_7 ./reentrantcorner.dgf
```

The results for the triangular grid shown in Figures 6, 7 and Table 2 may be reproduced from running

```
make poisson_alugrid_simplex_7
./poisson_alugrid_simplex_7 ./reentrantcorner.dgf
```

If an MPI implementation was found during the build process the tests may be run in parallel in an analogous way.

B Advanced user API

In this final appendix we want to describe two advanced aspects of the `dune-fem-hpdg` module, the definition of custom data projections for the restriction and prolongation of user data and the implementation of new discontinuous finite element spaces based on user-defined local basis function sets.

[3] `https://gitlab.dune-project.org/dune-fem/dune-fem`

[4] `https://gitlab.dune-project.org/christoph.gersbacher/dune-fem-hpdg.git`

[5] `http://www.dune-project.org/doc/installation-notes.html`

B.1 Restriction and prolongation of custom user data

Remember from Section 4.3 that the restriction and prolongation of data in `dune-fem-hpdg` is handled by the `hpDG::AdaptationManager` class. Any input must be encapsulated in a `DataProjection` object.

```
using DataProjection = hpDG::DefaultDataProjection<Data>;
hpDG::AdaptationManager<DiscreteFunctionSpace, DataProjection>
    p_AdaptManager(
    space, DataProjection(u));
```

A `DataProjection` is a callable object encapsulating the data $u^{(m)}$ associated with the discrete function space $X^{(m)}$. Any such projection must inherit from the base class `hpDG::DataProjection`, as illustrated by the following listing.

```
template <class DiscreteFunctionSpace, class Data>
class DefaultDataProjection
    : public DataProjection<DiscreteFunctionSpace,
                            DefaultDataProjection<Data> > {
 public:
  using BasisFunctionSetType =
      typename DiscreteFunctionSpace::BasisFunctionSetType;
  using EntityType = typename BasisFunctionSetType::EntityType;

  void DataProjection::operator()(const EntityType &element,
                                  const BasisFunctionSetType &former,
                                  const BasisFunctionSetType &future,
                                  const vector<size_t> &origin,
                                  const vector<size_t> &destination);
};
```

The method `operator()` is expected to evaluate the local projection $\Pi^{(m+1)} u^{(m)}\restriction_E$ for each element $E \in \mathcal{G}^{(m)} \cap \mathcal{G}^{(m+)1}$. Its input arguments are

i) the grid element E,

ii) the former and future local basis function sets $\left(\varphi_{E,i}^{(m)}\right)_{i=0,\dots,n_E^{(m)}-1}$, $\left(\varphi_{E,i}^{(m+1)}\right)_{i=0,\dots,n_E^{(m+1)}-1}$,

iii) and the DOF indices $(\mu^{(m)}(E,i))_{i=0,\dots,n_E^{(m)}-1}$, $(\mu^{(m+1/2)}(E,i))_{i=0,\dots,n_E^{(m+1)}-1}$.

Note that this is all information needed in order to implement the local L^2-projection from Example 3.1.

B.2 Setup of new adaptive discontinuous finite element spaces

As described in Section 4.2 the discrete function spaces provided by `dune-fem-hpdg` inherit from a common base class implementing almost all functionality, see Figure 3. In the same way, users may setup new discontinuous finite element spaces. To do so, users need to provide an implementation of the `hpDG::BasisFunctionSets` interface which contains essentially two methods.

```
template <class BasisFunctionSet, class Key>
class BasisFunctionSets {
 public:
  using EntityType = typename BasisFunctionSet::EntityType;

  BasisFunctionSet basisFunctionSet(const EntityType &entity,
                                    const Key &key) const;
  size_t blocks(GeometryType type, const Key &key) const;
```

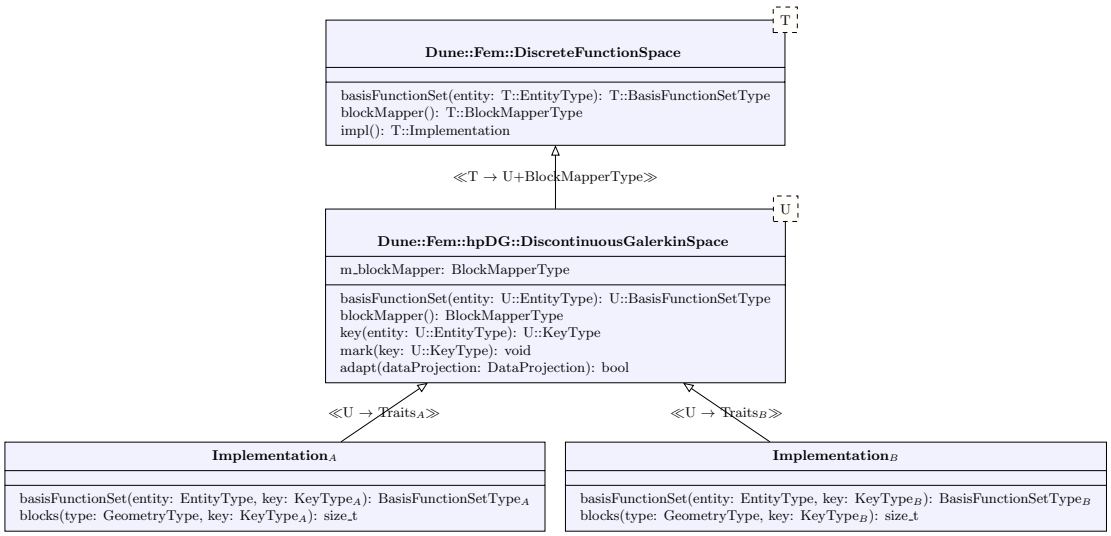

Figure 3: Class diagram for derived implementations of *hp*-adaptive discontinuous finite element spaces inheriting from the `hpDG::DiscontinuousGalerkinSpace` base class.

```
9     // ...
10  };
```

Remember that a `GeometryType` is a DUNE data type that identifies a reference element. Given a grid element E with reference element R_E and a key k the method `basisFunctionSet` returns the corresponding local basis function set $\mathcal{B}_{E,k}$. We assume that the size of the local basis function set only depends on R_E and k. Then, the method `blocks` returns the number of blocks to be reserved for $\mathcal{B}_{E,k}$. A `BlockMapper` is a DUNE-FEM class representation a DOF mapping; in case of a scalar function spaces the number of blocks coincides with the number of local basis functions in $\mathcal{B}_{E,k}$.

Users may inherit from the base implementation `hpDG::DiscontinuousGalerkinSpace`. Alternatively, the convenience function `make_space` immediately constructs a discrete function space, given a grid part, the family of basis function sets, and a default key for initializing the space.

```
1  template <class GridPart, class BasisFunctionSets>
2  std::unique_ptr<DefaultDiscontinuousGalerkinSpace<BasisFunctionSets> >
3  make_space(GridPart &gridPart, const BasisFunctionSets
       &basisFunctionSets,
4             const typename BasisFunctionSets::KeyType &key);
```

References

[1] M. Alkämper, A. Dedner, R. Klöfkorn, and M. Nolte. The DUNE-ALUGRID module. *Arch. Numer. Softw.*, 4(1):1–28, 2016.

[2] I. Babuška, B. A. Szabo, and I. N. Katz. The *p*-version of the finite element method. *SIAM J. Numer. Anal.*, 18(3):515–545, 1981.

[3] W. Bangerth, R. Hartmann, and G. Kanschat. deal.II – A general purpose object oriented finite element library. *ACM Trans. Math. Softw.*, 33(4):24/1–24/27, 2007.

[4] W. Bangerth and O. Kayser-Herold. Data structures and requirements for *hp* finite element software. *ACM Trans. Math. Softw.*, 36(1):4:1–4:31, 2009.

[5] P. Bastian, M. Blatt, A. Dedner, Ch. Engwer, R. Klöfkorn, R. Kornhuber, M. Ohlberger, and O. Sander. A generic grid interface for parallel and adaptive scientific computing. Part II: Implementation and tests in DUNE. *Computing*, 82(2–3):121–138, 2008.

[6] P. Bastian, M. Blatt, A. Dedner, Ch. Engwer, R. Klöfkorn, M. Ohlberger, and O. Sander. A generic grid interface for parallel and adaptive scientific computing. Part I: Abstract framework. *Computing*, 82(2–3):103–119, 2008.

[7] M. Blatt and P. Bastian. The iterative solver template library. In B. Kågström, E. Elmroth, J. Dongarra, and J. Waśniewski, editors, *Applied Parallel Computing. State of the Art in Scientific Computing*, pages 666–675. Springer Berlin Heidelberg, 2007.

[8] M. Blatt and P. Bastian. On the generic parallelisation of iterative solvers for the finite element method. *Int. J. Comput. Sci. Eng.*, 4(1):56–69, 2008.

[9] A. Dedner, R. Klöfkorn, and M. Kränkel. Continuous finite-elements on non-conforming grids using discontinuous Galerkin stabilization. In J. Fuhrmann, M. Ohlberger, and Ch. Rohde, editors, *Finite Volumes for Complex Applications VII. Methods and Theoretical Aspects*, pages 207–215. Springer International Publishing, 2014.

[10] A. Dedner, R. Klöfkorn, M. Nolte, and M. Ohlberger. A generic interface for parallel and adaptive discretization schemes: Abstraction principles and the DUNE-FEM module. *Computing*, 90(3–4):165–196, 2010.

[11] L. Demkowicz, W. Rachowicz, and Ph. Devloo. A fully automatic *hp*-adaptivity. *J. Sci. Comput.*, 17(1):117–142, 2002.

[12] V. Dolejší. *hp*-DGFEM for nonlinear convection-diffusion problems. *Math. Comput. Simulation*, 87:87–118, 2013.

[13] V. Dolejší and M. Feistauer. *Discontinuous Galerkin Method. Analysis and Applications to Compressible Flow.* Springer International Publishing, 2015.

[14] T. Eibner and J. Melenk. An adaptive strategy for *hp*-FEM based on testing for analyticity. *Comput. Mech.*, 39(5):575–595, 2007.

[15] Ph. Frauenfelder and Ch. Lage. Concepts - An object-oriented software package for partial differential equations. *ESAIM Math. Model. Numer. Anal.*, 36(5):937–951, 2002.

[16] W. Gui and I. Babuška. The h, p and h-p versions of the finite element method in one dimension. Part I. The error analysis of the p-version. *Numer. Math.*, 49(6):577–612, 1986.

[17] W. Gui and I. Babuška. The h, p and h-p versions of the finite element method in one dimension. Part II. The error analysis of the h- and h-p versions. *Numer. Math.*, 49(6):613–657, 1986.

[18] W. Gui and I. Babuška. The h, p and h-p versions of the finite element method in one dimension. Part III. The adaptive h-p-version. *Numer. Math.*, 49(6):659–683, 1986.

[19] B. Guo and I. Babuška. The h-p version of the finite element method. Part I. The basic approximation results. *Comput. Mech.*, 1(1):21–41, 1986.

[20] B. Guo and I. Babuška. The h-p version of the finite element method. Part II. General results and applications. *Comput. Mech.*, 1(3):203–220, 1986.

[21] P. Houston, D. Schötzau, and Th. P. Wihler. Energy norm a posteriori error estimation of hp-adaptive discontinuous Galerkin methods for elliptic problems. *Math. Models Methods Appl. Sci.*, 17(01):33–62, 2007.

[22] P. Houston, B. Senior, and E. Süli. *Numerical Mathematics and Advanced Applications*, chapter Sobolev regularity estimation for hp-adaptive finite element methods, pages 631–656. Springer Milan, 2003.

[23] P. Houston and E. Süli. A note on the design of hp-adaptive finite element methods for elliptic partial differential equations. *Comput. Method. Appl. M.*, 194(2–5):229–243, 2005.

[24] P. Houston, E. Süli, and Th. P. Wihler. A posteriori error analysis of hp-version discontinuous Galerkin finite-element methods for second-order quasi-linear elliptic PDEs. *IMA J. Numer. Anal.*, 28(2):245–273, 2008.

[25] R. Klöfkorn. *Numerics for evolution equations: A general interface based design concept.* Doctoral dissertation, Albert-Ludwigs-Universität Freiburg, 2009.

[26] C. Lage. Concept oriented design of numerical software. Research Report 98-07, Seminar for Applied Mathematics, ETH Zürich, 1998.

[27] W. F. Mitchell. PHAML user's guide. Technical Report NISTIR 7374, National Institute of Standards and Technology, 2006.

[28] W. F. Mitchell and M. A. McClain. A comparison of hp-adaptive strategies for elliptic partial differential equations (long version). Technical Report 7824, National Institute of Standards and Technology, 2011.

[29] W. F. Mitchell and M. A. McClain. A comparison of hp-adaptive strategies for elliptic partial differential equations. *ACM Trans. Math. Softw.*, 41(1):2:1–2:39, 2014.

[30] B. Schupp. *Entwicklung eines effizienten Verfahrens zur Simulation kompressibler Strömungen in 3D auf Parallelrechnern.* Doctoral dissertation, Albert-Ludwigs-Universität Freiburg, 1999.

[31] P. Šolín and L. Demkowicz. Goal-oriented hp-adaptivity for elliptic problems. *Comput. Method. Appl. M.*, 193(6–8):449–468, 2004.

[32] L. Zhu, S. Giani, P. Houston, and D. Schötzau. Energy norm a posteriori error estimation for hp-adaptive discontinuous Galerkin methods for elliptic problems in three dimensions. *Math. Models Methods Appl. Sci.*, 21(02):267–306, 2011.

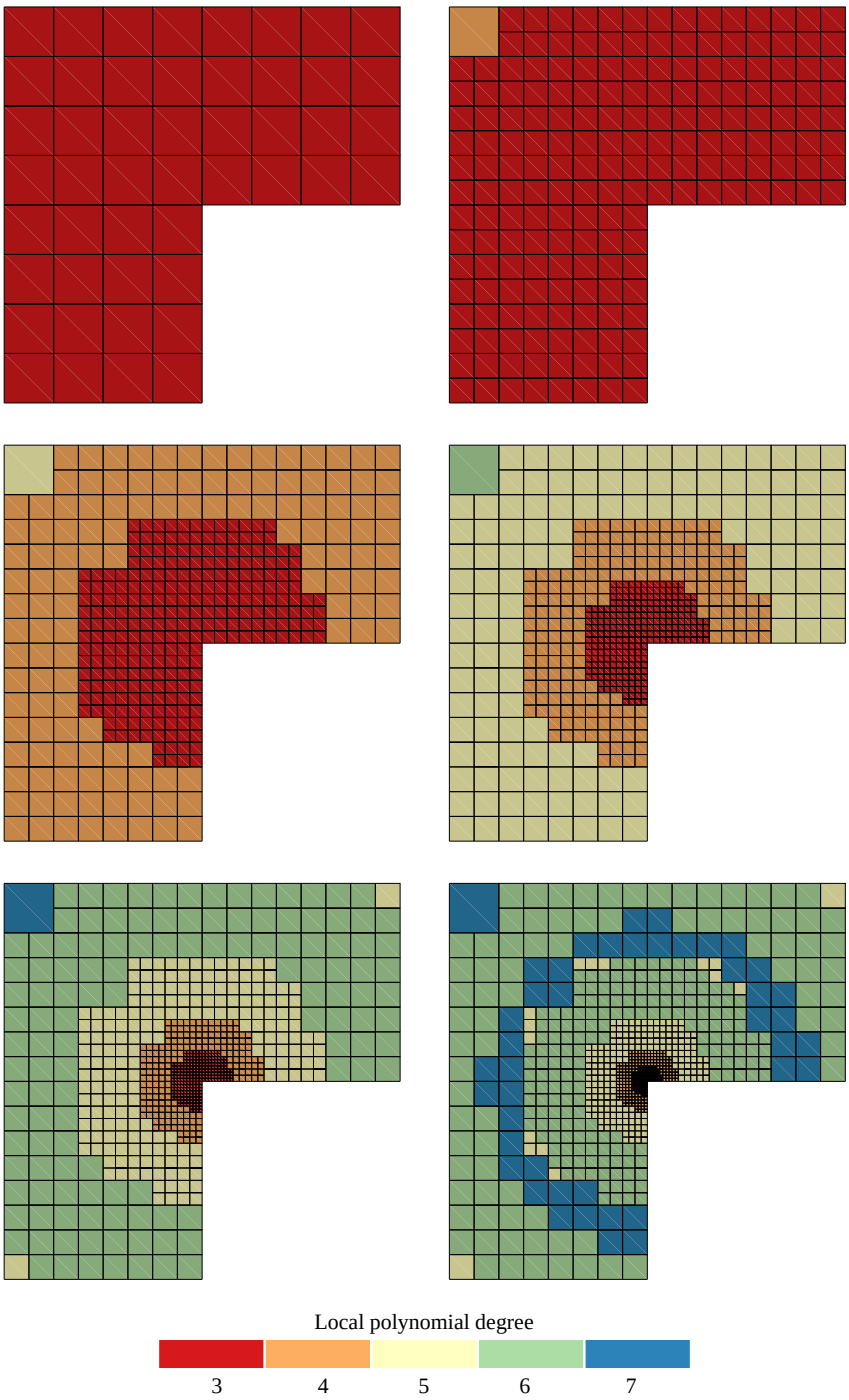

Figure 4: *hp*-meshes generated by the *hp*-adaptive SIPG method when solving the reentrant corner benchmark problem on a quadrilateral grid.

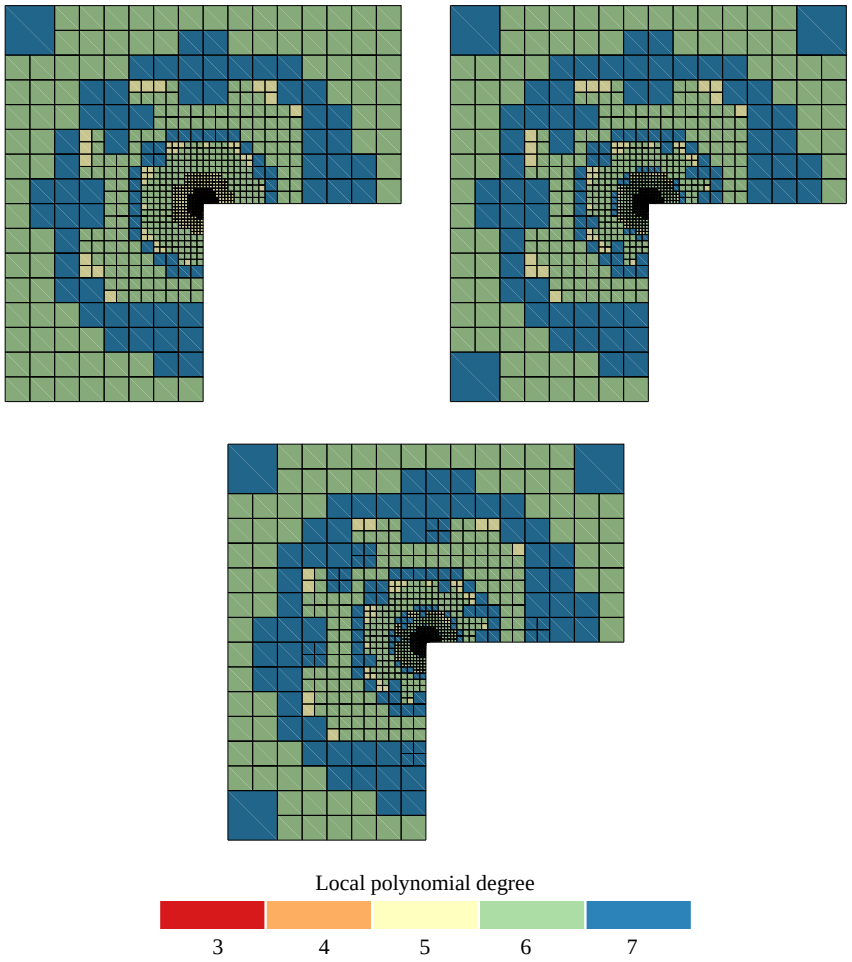

Figure 5: *hp*-meshes generated by the *hp*-adaptive SIPG method when solving the reentrant corner benchmark problem on a quadrilateral grid.

Table 1: Number of elements and degrees of freedom, errors and convergence rates, and effectivity index for the series of quadrilateral grids depicted in Figures 4 and 5.

Elements	DoFs	$\|u-u_h\|_{L^2}$	EOC	$\|u-u_h\|_{DG}$	EOC	$\left(\sum \eta_E^2\right)^{\frac{1}{2}}$	Eff. index
48	480	2.30×10^{-3}	—	2.22×10^{-1}	—	3.30×10^{-1}	1.49
189	1895	8.56×10^{-4}	1.4	1.39×10^{-1}	0.7	2.08×10^{-1}	1.49
378	4416	3.25×10^{-4}	2.3	8.77×10^{-2}	1.1	1.31×10^{-1}	1.49
567	8008	1.25×10^{-4}	3.2	5.52×10^{-2}	1.6	8.23×10^{-2}	1.49
756	12 846	4.85×10^{-5}	4.0	3.48×10^{-2}	2.0	5.19×10^{-2}	1.49
945	18 376	1.90×10^{-5}	5.2	2.19×10^{-2}	2.6	3.27×10^{-2}	1.49
1104	22 964	7.47×10^{-6}	8.4	1.38×10^{-2}	4.1	2.06×10^{-2}	1.49
1269	27 480	2.95×10^{-6}	10.4	8.70×10^{-3}	5.1	1.30×10^{-2}	1.49
1464	32 712	1.17×10^{-6}	10.6	5.48×10^{-3}	5.3	8.17×10^{-3}	1.49

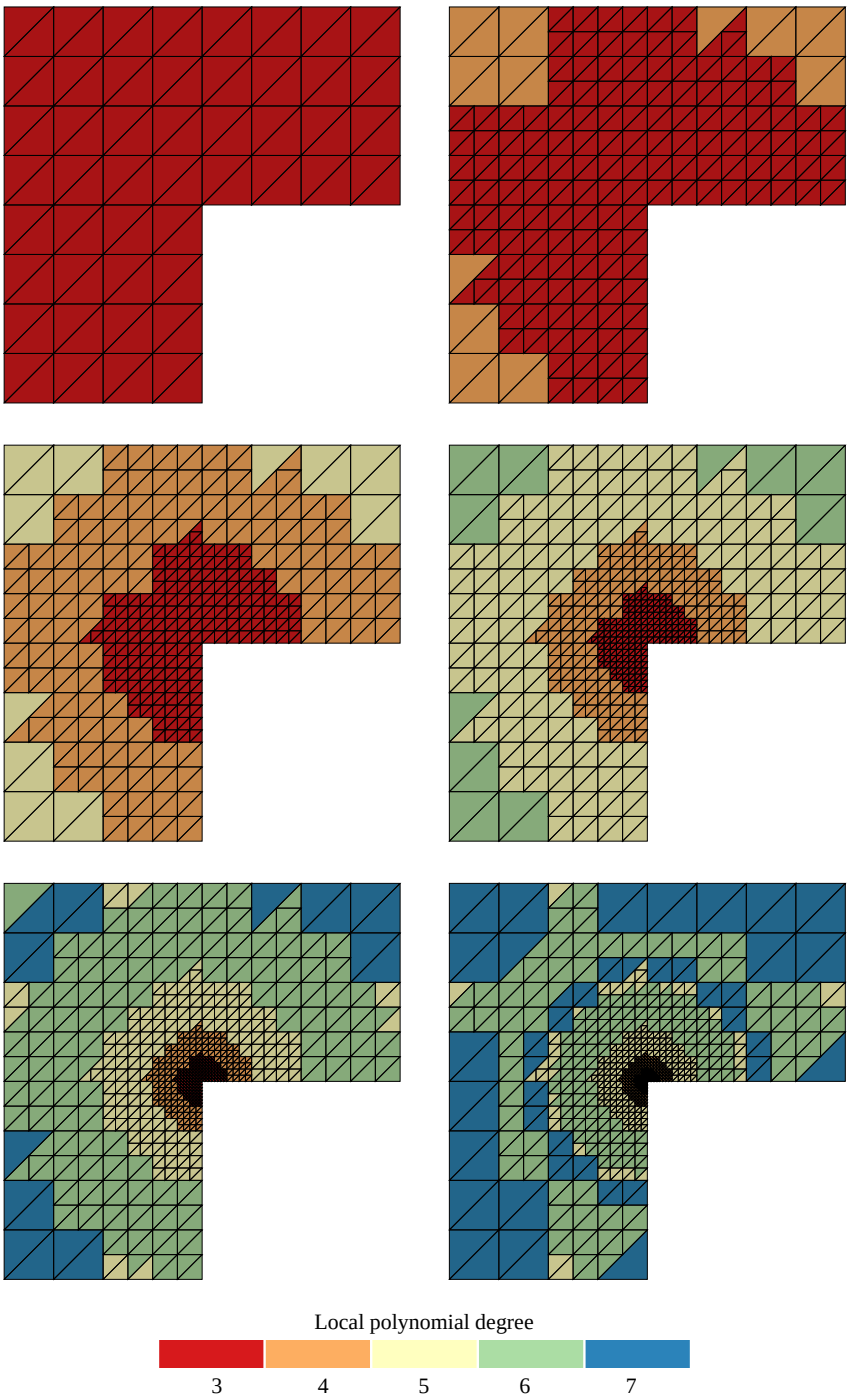

Figure 6: *hp*-meshes generated by the *hp*-adaptive SIPG method when solving the reentrant corner benchmark problem on an triangular grid.

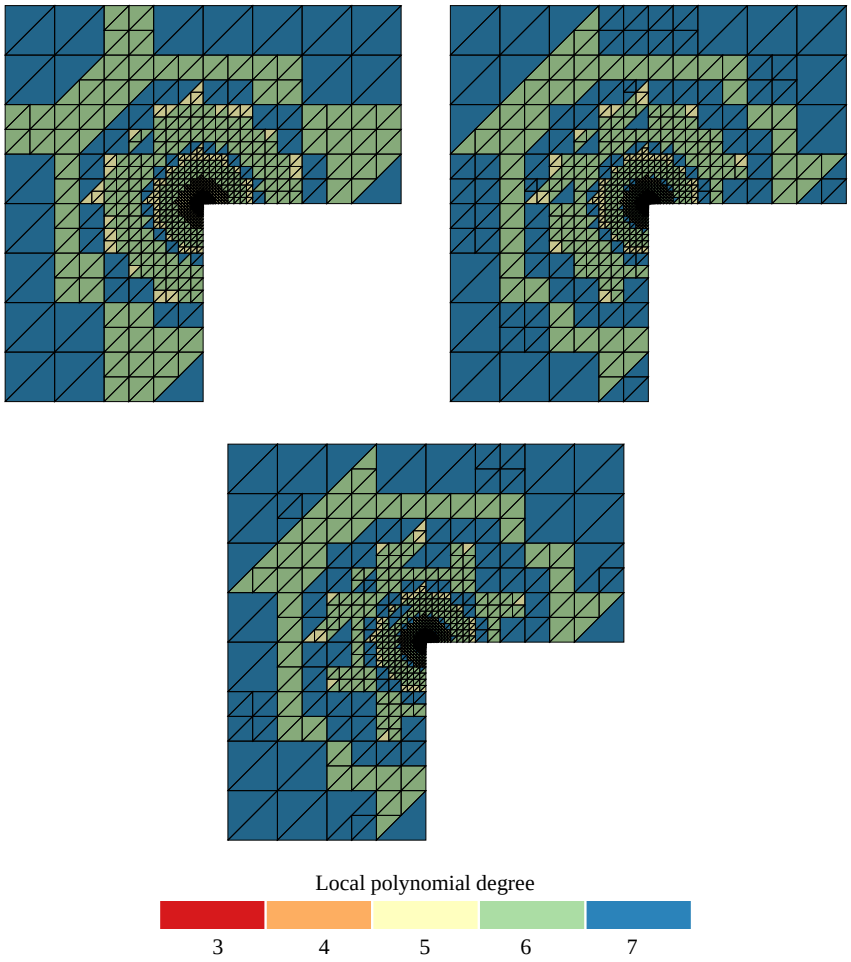

Figure 7: *hp*-meshes generated by the *hp*-adaptive SIPG method when solving the reentrant corner benchmark problem on an triangular grid.

Table 2: Number of elements and degrees of freedom, errors and convergence rates, and effectivity index for the series of triangular grids depicted in Figures 6 and 7.

Elements	DoFs	$\|u-u_h\|_{L^2}$	EOC	$\|u-u_h\|_{DG}$	EOC	$\left(\sum \eta_E^2\right)^{\frac{1}{2}}$	Eff. index
96	960	1.08×10^{-3}	—	2.31×10^{-1}	—	3.04×10^{-1}	1.32
318	3290	4.09×10^{-4}	1.6	1.44×10^{-1}	0.8	1.92×10^{-1}	1.33
558	6930	1.57×10^{-4}	2.6	9.06×10^{-2}	1.2	1.21×10^{-1}	1.33
798	12 014	6.08×10^{-5}	3.4	5.70×10^{-2}	1.7	7.60×10^{-2}	1.33
1038	18 690	2.37×10^{-5}	4.3	3.59×10^{-2}	2.1	4.79×10^{-2}	1.33
1227	24 137	9.33×10^{-6}	7.3	2.26×10^{-2}	3.6	3.02×10^{-2}	1.34
1464	30 643	3.68×10^{-6}	7.8	1.42×10^{-2}	3.9	1.90×10^{-2}	1.34
1716	37 871	1.46×10^{-6}	8.8	8.97×10^{-3}	4.4	1.20×10^{-2}	1.34
1899	42 361	5.76×10^{-7}	16.5	5.65×10^{-3}	8.2	7.54×10^{-3}	1.34